

2MW
LIQUID COOLED
CONTAINERISED
DATA CENTER




**PUE
1.03**
3%
INEFFICIENCY OF
INFRASTRUCTURE




**UP TO
120kW**
PER RACK

16x
200W NODES
PER 1u



1.6Tb/s
TRIFECTA
INTERCONNECTS



16x
COM Express
SERVERS
PER 1u

1u
INCLUDES
HIGH SPEED &
MGT SWITCHES

European co-design for exascale applications

The Horizon 2020 EuroEXA project proposes a ground-breaking design for mind-blowing results: Over four times higher performance and four times higher energy efficiency than today's High-Performance platforms. The goal of the project is develop a Xilinx FPGA board running at several TFlop/s, liquid cool and connect a few hundred and interconnect them to create a few PFlop/s cluster.

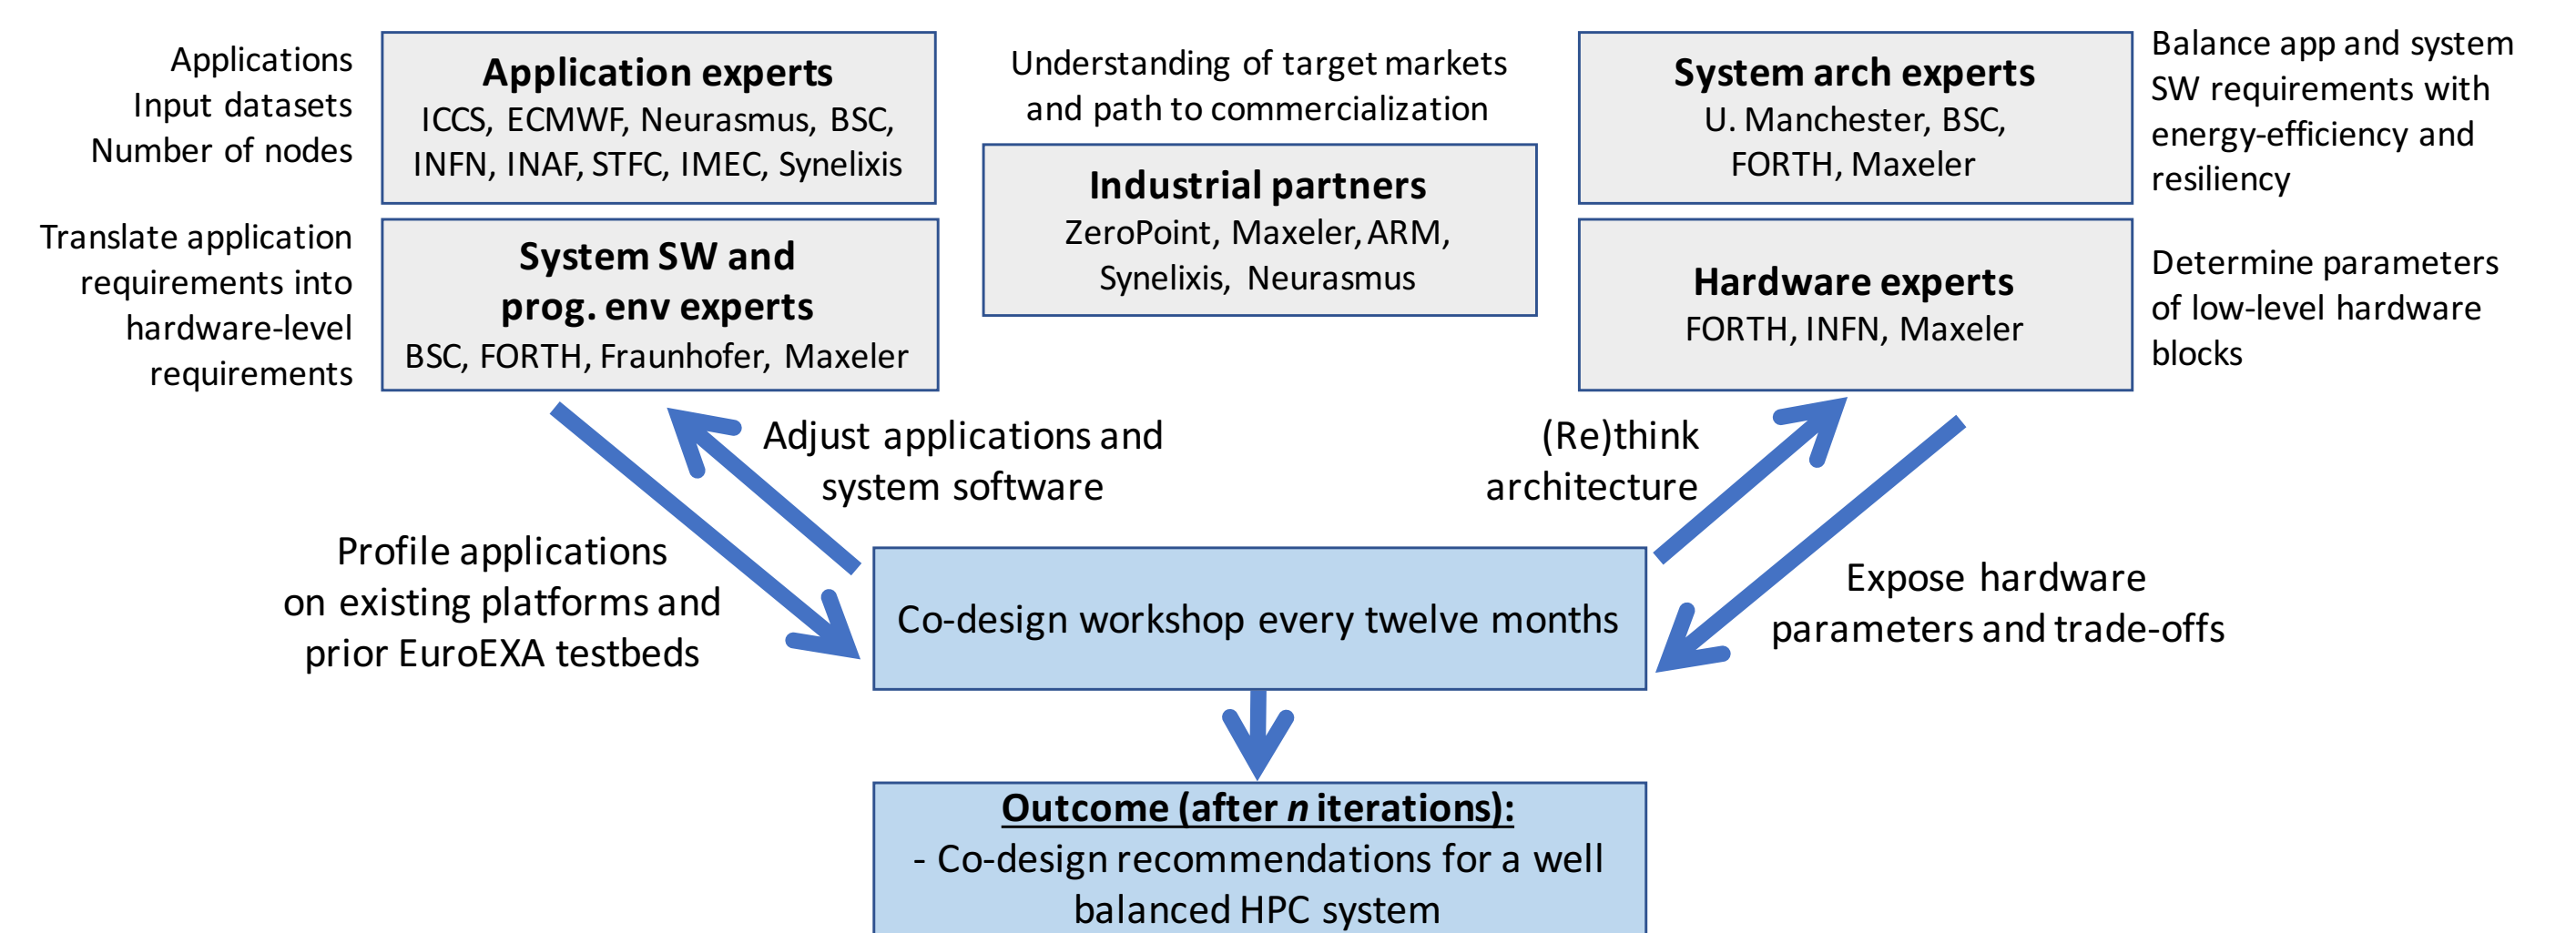
Vision

- First testbed architecture will be shown to be capable of scaling to world-class peak performance in excess of 400 PFLOPS with an estimated system power of around 30 MW peak.
- A compute-centric 250 PFLOPS per 15 MW by 2019.
- Show that an exascale machine could be built in 2020 within 30 shipping containers with an edge to edge distance of less than 40m.

Who are the partners?

<p>Institute of Communication and Computer Systems ARM Fraunhofer Iceotope Technologies Limited IMEC VZW Maxeler Neurasmus BV Synelxis Solutions S.A</p>	<p>The Barcelona Supercomputing Center The European Centre for Medium-Range Weather Forecasts The Foundation for Research and Technology – HELLAS The Italian National Institute for Astrophysics The Italian National Institute for Nuclear Physics The Science and Technology Facilities Council University of Manchester ZeroPoint Technologies</p>
--	--

Architecture

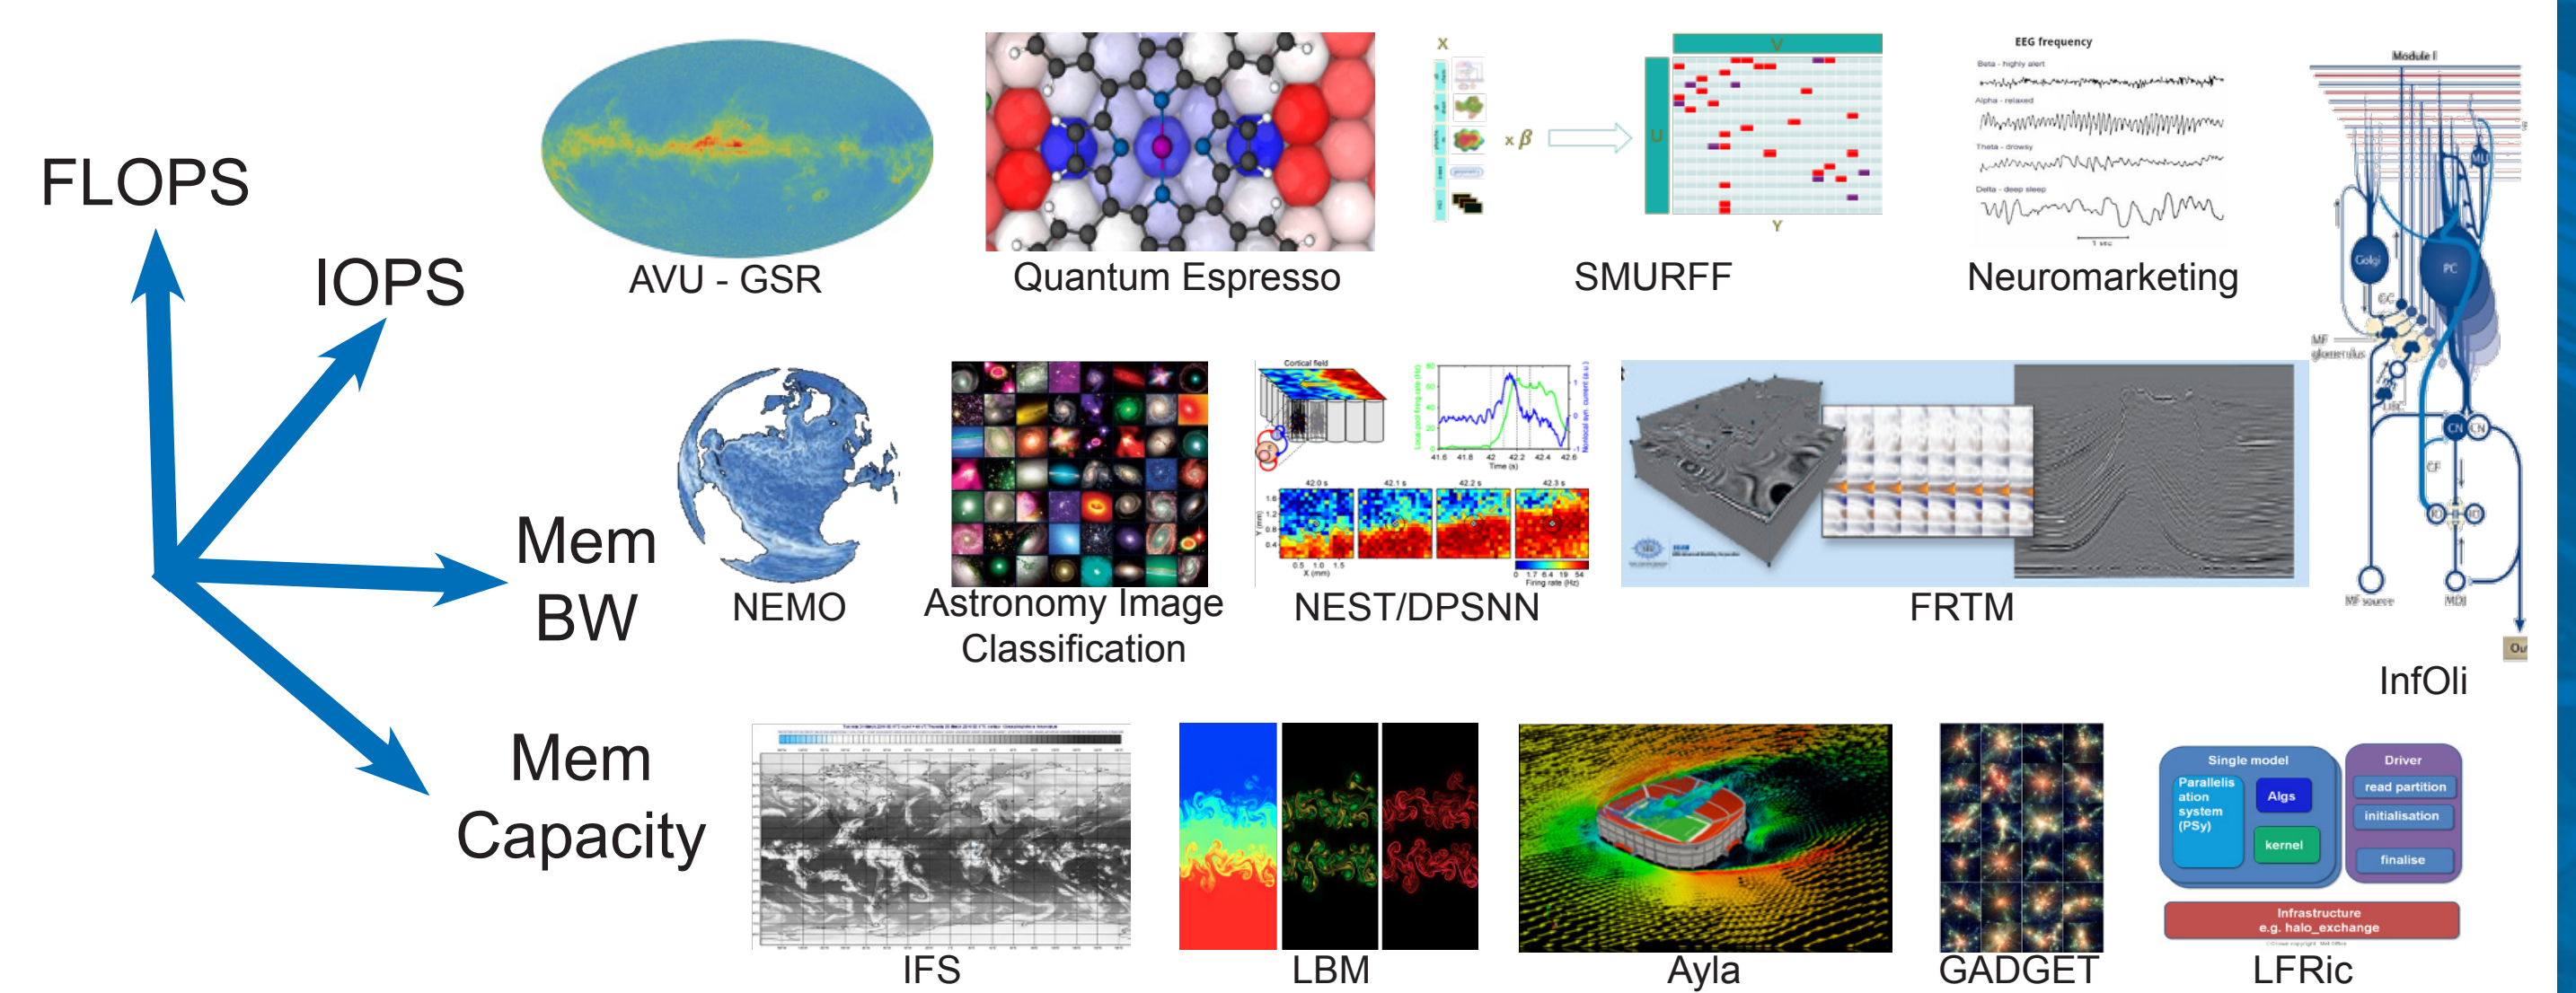


Trifecta Scalable I/O

As node count increases, homogeneous interconnects suffer untenable pressure on the per-node interconnect radix or suffer from multi-hop latencies, longer cables, and increased communication energy requirements. EuroEXA uses the Trifecta Scalable Interconnect, a proximity optimized interconnect providing data and communication locality between various levels of hierarchy of peers.


- **T1** Direct Node to Node Interconnect in groups of 4 Nodes – 320Gb/s Per Node
- **T2** Node to Blade Level FPGA Switch, in groups of 16 Nodes – 64Gb/s Per Node
- **T3** Cabinet to Cabinet Over Commodity Switches, Using SDN Openflow – 800Gb/s Per Blade

Applications



CRDB - FPGA heavy nodes

- Co-Design Influenced
- VU9P Xilinx FPGAs
- ARM CPU and FPGA for Networking and Storage
- 64GB RAM
- 500GB SSD
- 4x16Gb T2 Trifecta I/O
- 20x16Gb T1 Trifecta I/O



Scale out testbed

- Deployed Q3&Q4 2020
- 256 CRDB FPGA Heavy Nodes
- 256 New Processor Heavy Nodes
- 48V DC Distribution
- Shipping Container Data Centre Design
- Total Liquid Cooling
- Hot Water Out, Chillerless Design

